arXiv:1609.08913v1 [stat.ML] 28 Sep 2016

The Famine of Forte: Few Search Problems Greatly Favor Your Algorithm

George D. Montañez
Machine Learning Department
Carnegie Mellon University
Pittsburgh, Pennsylvania, USA
gmontane@cs.cmu.edu

Abstract

No Free Lunch theorems show that the average performance across any closed-under-permutation set of problems is fixed for all algorithms, under appropriate conditions. Extending these results, we demonstrate that the proportion of favorable problems is itself strictly bounded, such that no single algorithm can perform well over a large fraction of possible problems. Our results explain why we must either continue to develop new learning methods year after year or move towards highly parameterized models that are both flexible and sensitive to their hyperparameters.

1 Introduction

BIGCITY, United States. In a fictional world not unlike our own, a sophisticated criminal organization plots to attack an unspecified landmark within the city. Due to the complexity of the attack and methods of infiltration, the group is forced to construct a plan relying on the coordinated actions of several interdependent agents, of which the failure of any one would cause the collapse of the entire plot. As a member of the city's security team, you must allocate finite resources to protect the many important locations within the city. Although you know the attack is imminent, your sources have not indicated which location, of the hundreds possible, will be hit; your lack of manpower forces you to make assumptions about target likelihoods. You know you can foil the plot if you stop even one enemy agent. Because of this, you seek to maximize the odds of capturing an agent by placing vigilant security forces at the strategic locations throughout the city. Allocating more security to a given location increases surveillance there, raising the probability a conspirator will be found if operating nearby. Unsure of your decisions, you allocate based on your best information, but continue to second-guess yourself.

With this picture in mind, we can analyze the scenario through the lens of algorithmic search. Algorithmic search methods, whether employed in fictional security situations or by researchers in the lab, all share common elements. They have a **search space** of elements (possible locations) which contains items of high value at unknown locations within that space. These high-value elements are **targets** of the search, which are sought by a process of inspecting elements of the search space to see if they contain any elements of the target set. (We refer to the inspection of search space elements as **sampling** from the search space, and each sampling action as a **query**.) In our fictional scenario, the enemy activity locations constitute the target set, and surveillance at city locations corresponds to a search process (i.e., attempting to locate and arrest enemy agents). The search process can be deterministic or contain some degree of randomization, such as choosing which elements to inspect using a weighted distribution over the search space. The search process is a **success** if an element of the target set is located during the course of inspecting elements of the space. The **history** of the search refers to the collection of elements already sampled and any accompanying information (e.g., the intelligence data gathered thus far by the security forces).

There are many ways to search such a space. Research on algorithmic search is replete with proposed search methods that attempt to increase the expected likelihood of search success, across many different problem domains. However, a direct result of the No Free Lunch theorems are that no search method is universally optimal across all search problems [23, 3, 8, 17, 22, 10, 7, 6, 4], since all algorithmic search methods have performance equivalent to uniform random sam-

pling on the search space when uniformly averaged across any closed-under-permutation set of problem functions [17]. Thus, search (and learning) algorithms must trade weaker performance on some problems for improved performance on others [23, 16]. Given the fact that there exists no universally superior search method, we typically seek to develop methods that perform well over a range of important problems. An important question naturally arises:

> *Is there a limit to the number of problems for which a search algorithm can perform well?*

For search problems in general, the answer is *yes*, as the proportion of search problems for which an algorithm outperforms uniform random sampling is strictly bounded in relation to the degree of performance improvement. Previous results have relied on the fact that most search problems have many target elements, and thus are not difficult for random sampling [15, 9]. Since the typical search problem is not difficult for random sampling, it becomes hard to find a problem for which an algorithm can significantly and reliably outperform random sampling. While theoretically interesting, a more relevant situation occurs when we consider difficult search problems, which are those having relatively small target sets. We denote such problems as **target-sparse problems** or following [15], **target-sparse functions**. Note, this use of sparseness differs from that typical in the machine learning literature, since it refers to sparseness of the target space, not sparseness in the feature space. The focus of this paper is to prove results that continue to hold even for target-sparse functions.

2 Contributions

We show that favorable search problems must necessarily be rare, regardless of target-sparseness. We extend horizontal No Free Lunch type results (namely, that the mean performance across sets of problems is fixed for all algorithms) to show that the proportion of favorable problems is also strictly bounded in relation to the inherent problem difficulty and the degree of improvement sought (i.e., not just the mean performance is bounded). Our work represents a breakthrough in this regard.

Next, by using active information to measure performance improvement in bits [5], we prove a conservation of information result [10, 5, 15], showing that the amount of information required to locate a search problem giving an expected b bits of information is at least b bits.

We then prove a search-for-a-search result [5, 4], showing that favorable search strategies are just as rare as favorable search problems.

We relate the degree of dependence (measured in mutual information) between target sets and external information resources, such as objective functions, noisy measurements[1] or sets of training data, to the maximum improvement in search performance. We prove that for a fixed target-sparseness and given an algorithm with expected per-query probability of success q,

$$q \leq \frac{I(T;F) + D(P_T \| \mathcal{U}_T) + 1}{I_\Omega} \tag{1}$$

where $I(T;F)$ is the mutual information between target set T (as a random variable) and external information resource F, $D(P_T \| \mathcal{U}_T)$ is the Kullback-Liebler divergence between the marginal distribution on T and the uniform distribution on target sets, and I_Ω is the baseline information cost for the search problem due to sparseness. This surprisingly simple equation takes into account degree of dependence, target sparseness, target function uncertainty, and the contribution of random luck.

Lastly, we apply the results to several real-world problem settings, showing how these results lead to new insights in different research areas.

3 Framework

We begin by formalizing our problem setting, search method abstraction and other necessary concepts.

3.1 The Search Problem

We abstract the search problem as follows. The search space, denoted Ω, contains the elements to be examined. We limit ourselves to finite, discrete search spaces, which entails no loss of generality when considering search spaces fully representable on physical computer hardware within a finite time. Let the **target set** $T \subset \Omega$ be a nonempty subset of the search space. The set T can be represented using a binary vector the size of Ω, where an indexed position evaluates to 1 whenever the corresponding element is in T and 0 otherwise. Thus, each T corresponds to exactly one binary vector of length $|\Omega|$, and vice versa. We refer to this one-to-one mapped binary vector as a **target function** and use the terms target function and target set interchangeably, depending on context. These target sets/functions will help us define our space of possible search problems, as we will see shortly.

[1] Noisy in the sense of inaccurate or biased, not in the sense of indeterministic.

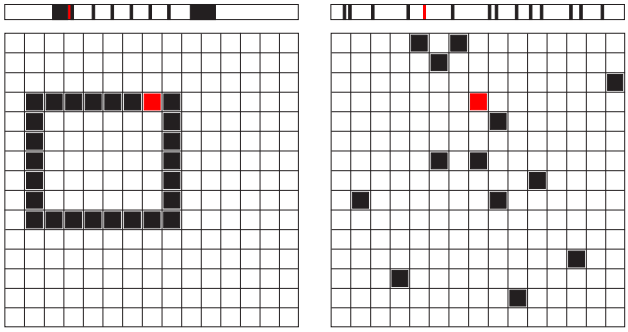

Figure 1: Two example target sets. In one case (left), knowing the location of the black target elements fully determines the location of the red target element. In the second case (right), the locations of the black target elements reveal nothing concerning the location of the red element. Problems are represented as a binary vector (top row), and two-dimensional search space (bottom row).

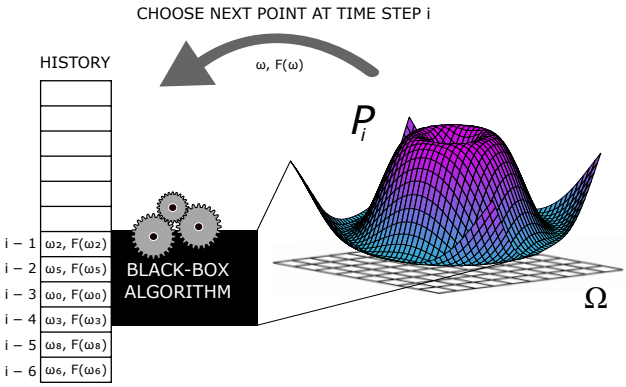

Figure 2: Black-box search algorithm. At time i the algorithm computes a probability distribution P_i over the search space Ω, using information from the history, and a new point is drawn according to P_i. The point is evaluated using external information resource F. The tuple $(\omega, F(\omega))$ is then added to the history at position i.

Figure 1 shows two example target sets, in binary vector and generic target set form. The set on the left has strong (potentially exploitable) structure governing the placement of target elements, whereas the example on the right is more or less random. Thus, knowing the location of some target elements may or may not be able to help one find additional elements, and in general there may be any degree of correlation between the location of target elements already uncovered and those yet to be discovered.

Typically, elements from the search space Ω are evaluated according to some **external information resource**, such as an objective function f. We abstract this resource as simply a finite length bit string, which could represent an objective function, a set of training data (in supervised learning settings), or anything else. The resource can exist in coded form, and we make no assumption about the shape of this resource or its encoding. Our only requirement is that it can be used as an oracle, given an element (possibly the null element) from Ω. In other words, we require the existence of two methods, one for initialization (given no element) and one for evaluating queried points in Ω. We assume both methods are fully specified by the problem domain and external information resource. With a slight abuse of notation, we define an **information resource evaluation** as $F(\omega) := g(F, \omega)$, where g is an extraction function applied to the information resource and $\omega \in \Omega \cup \{\emptyset\}$. Therefore, $F(\emptyset)$ represents the method used to extract initial information for the search algorithm (absent of any query), and $F(\omega)$ represents the evaluation of point ω under resource F.

A **search problem** is defined as a 3-tuple, (Ω, T, F), consisting of a search space, a target subset of the space and an external information resource F, respectively. Since the target locations are hidden, any information gained by the search concerning the target locations is mediated through the external information resource F alone. Thus, the space of possible search problems includes many deceptive search problems, where the external resource provides misleading information about target locations, similar to when security forces are given false intelligence data, and many noisy problems, similar to when imprecise intelligence gathering techniques are used. In the fully general case, there can be any relationship between T and F. Because we consider any and all degrees of dependence between external information resources and target locations, this effectively creates independence when considering the set as a whole, allowing the first main result of this paper to follow as a consequence.

However, in many natural settings, target locations and external information resources are tightly coupled. For example, we typically threshold objective function values to designate the target elements as those that meet or exceed some minimum. Doing so enforces dependence between target locations and objective functions, where the former is fully determined by the latter, once the threshold is known. This dependence causes direct correlation between the objective function (which is observable) and the target locations (which are not directly observable). We will demonstrate such correlation is exploitable, affecting the upper bound on the expected probability of success.

3.2 The Search Algorithm

An algorithmic search is any process that chooses elements of a search space to examine, given some history of elements already examined and information resource evaluations. The history consists of two parts: a query trace and a resource evaluation trace. A **query trace** is a record of points queried, indexed by time. A **resource evaluation trace** is a record of partial information extracted from F, also indexed by time. The information resource evaluations can be used to build elaborate predictive models (as is the case of Bayesian optimization methods), or ignored completely (as is done with uniform random sampling). The algorithm's internal state is updated at each time step (according to the rules of the algorithm), as new query points are evaluated against the external information resource. The search is thus an iterative process, with the history h at time i represented as $h_i = (\omega, F(\omega))$. The problem domain defines how much initial information from F is given by $F(\emptyset)$, as well as how much (and what) information is extracted from F at each query.

We allow for both deterministic and randomized algorithms, since deterministic algorithms are equivalent to randomized algorithms with degenerate probability functions (i.e., they place all probability mass on a single point). Furthermore, any population based method can also be represented in this framework, by holding P fixed when selecting the m elements in the population, then considering the history (possibly limiting the horizon of the history to only m steps, creating a Markovian dependence on the previous population) and updating the probability distribution over the search space for the next m steps.

Abstracting away the details of how such a distribution P is chosen, we can treat the search algorithm as a black box that somehow chooses elements of a search space to evaluate. A search is successful if an element in the target set is located during the search. Algorithm 1 outlines the steps followed by the black-box search algorithm and Figure 2 visually demonstrates the process.

Algorithm 1 Black-box Search Algorithm

```
Initialize h_0 ← (∅, F(∅)).
for all i = 1, ..., i_max do
    Using history h_{0:i−1}, compute P_i, the distribution over Ω.
    Sample element ω according to P_i.
    Set h_i ← (ω, F(ω)).
end for
if an element of T is contained in any tuple of h then
    Return success.
else
    Return failure.
end if
```

Our framework is sufficiently general to apply to supervised learning, genetic algorithms, and sequential model-based hyperparameter optimization, yet is specific enough to allow for precise quantitative statements. Although faced with an inherent balancing act whenever formulating a new framework, we emphasize greater generality to allow for the largest number of domain specific applications.

4 Measuring Performance

Since general search algorithms may vary the total number of sampling steps performed, we measure performance using the expected per-query probability of success,

$$q(T, F) = \mathbb{E}_{\tilde{P},H}\left[\frac{1}{|\tilde{P}|}\sum_{i=1}^{|\tilde{P}|} P_i(\omega \in T) \;\middle|\; F\right], \quad (2)$$

where $\tilde{P}$ is the sequence of probability distributions for sampling elements (with one distribution P_i for each time step i), H is the search history, and the T and F make the dependence on the search problem explicit. $|\tilde{P}|$ denotes the length of the sequence $\tilde{P}$, which equals the number of queries taken. The expectation is taken over all sources of randomness, which includes randomness over possible search histories and any randomness in constructing the various P_i from $h_{0:i-1}$ (if such a construction is not entirely deterministic). Taking the expectation over all sources of randomness is equivalent to taking the expected performance for samples drawn from an appropriately averaged distribution $\overline{P}(\cdot \mid F)$ (see Lemma 1). Because we are sampling from a fixed (after conditioning and expectation) probability distribution, the expected per-query probability of success for the algorithm is equivalent to the induced amount of probability mass allocated to target elements. Revisiting our fictional security scenario, the probability of capturing an enemy agent is proportional to the amount of security at his location. In the same way, the expected per-query probability of success is equal to the amount of induced probability mass at locations where target elements are placed. Thus, each $\overline{P}(\cdot \mid F)$ demarks an equivalence class of search algorithms mapping to the same averaged distribution; we refer to these equivalence classes as **search strategies**.

We use uniform random sampling with replacement as our baseline search algorithm, which is a simple, always available strategy and we define $p(T, F)$ as the per-query probability of success for that method. In a natural way, $p(T, F)$ is a measure of the intrinsic difficulty of a search problem [5], absent of any side-information. The ratio $p(T, F)/q(T, F)$ quantifies the

improvement achieved by a search algorithm over simple random sampling. Like an error quantity, when this ratio is small (i.e., less than one) the performance of the algorithm is better than uniform random sampling, but when it is larger than one, performance is worse. We will often write $p(T, F)$ simply as p, when the target set and information resource are clear from the context.

5 Main Results

We state our main results here with brief discussion, and with proofs given in the Appendix.

5.1 Famine of Forte

Theorem 1. *(Famine of Forte) Define*

$$\tau_k = \{T \mid T \subset \Omega, |T| = k \in \mathbb{N}\}$$

and let $\mathcal{B}_m$ denote any set of binary strings, such that the strings are of length m or less. Let

$$R = \{(T, F) \mid T \in \tau_k, F \in \mathcal{B}_m\}, \text{ and}$$
$$R_{q_{min}} = \{(T, F) \mid T \in \tau_k, F \in \mathcal{B}_m, q(T, F) \geq q_{min}\},$$

where $q(T, F)$ is the expected per-query probability of success for algorithm $\mathcal{A}$ on problem (Ω, T, F). Then for any $m \in \mathbb{N}$,

$$\frac{|R_{q_{min}}|}{|R|} \leq \frac{p}{q_{min}}$$

and

$$\lim_{m\to\infty} \frac{|R_{q_{min}}|}{|R|} \leq \frac{p}{q_{min}}$$

where $p = k/|\Omega|$.

We see that for small p (problems with sparse target sets) the proportion of favorable search problems is rare if we desire a strong probability of success. The larger $q_{\min}$, the smaller the proportion. In many real-world settings, we are given a difficult search problem (with minuscule p) and we hope that our algorithm has a reasonable chance of achieving success within a limited number of queries. According to this result, the proportion of problems fulfilling such criteria is also minuscule. Only if we greatly relax the minimum performance demanded, so that $q_{\min}$ approaches the scale of p, can we hope to easily stumble upon such accommodating search problems. If not, insight and domain knowledge are required to find problems for which an algorithm excels.

This result has many applications, as will be shown in Section 6.

5.2 Conservation of Information

Corollary 1. *(Conservation of Expected Active Information) Define* active information *as $I_+ = -\log p/q$, where p is the per-query probability of success for uniform random sampling and q is the probability of success for an alternative search algorithm. Define*

$$\tau_k = \{T \mid T \subset \Omega, |T| = k \in \mathbb{N}\}$$

and let $\mathcal{B}_m$ denote any set of binary strings, such that the strings are of length m or less. Let

$$R = \{(T, F) \mid T \in \tau_k, F \in \mathcal{B}_m\}, \text{ and}$$
$$R_b = \{(T, F) \mid T \in \tau_k, F \in \mathcal{B}_m, \mathbb{E}[I_+] \geq b\}.$$

Then for any $m \in \mathbb{N}$

$$\frac{|R_b|}{|R|} \leq 2^{-b}.$$

Active information [5] quantifies improvement in a search over a baseline search method, measuring gains in search performance in terms of information (bits). This allows us to precisely quantify the proportion of sparse favorable functions in relation to the number b bits improvement desired. Active information provides a nice geometric interpretation of improved search performance, namely, that the improved search method is equivalent to a uniform random sampler on a reduced search space. Similarly, a degraded-performance search is equivalent to a uniform random sampler (with replacement) on a larger search space.

Finding a search problem for which an algorithm effectively reduces the search space by b bits requires at least b bits, so information is conserved in this context. Assuming you have no domain knowledge to guide the process of finding a search problem for which your search algorithm excels, you are unlikely to stumble upon one by chance; indeed, they are exponentially rare in the amount of improvement sought.

5.3 Famine of Favorable Strategies

Theorem 2. *(Famine of Favorable Strategies) For any fixed search problem (Ω, t, f), set of probability mass functions $\mathcal{P} = \{P : P \in \mathbb{R}^{|\Omega|}, \sum_j P_j = 1\}$, and a fixed threshold $q_{min} \in [0, 1]$,*

$$\frac{\mu(\mathcal{G}_{t,q_{min}})}{\mu(\mathcal{P})} \leq \frac{p}{q_{min}},$$

where $\mathcal{G}_{t,q_{min}} = \{P : P \in \mathcal{P}, t^\top P \geq q_{min}\}$ and μ is Lebesgue measure. Furthermore, the proportion of possible search strategies giving at least b bits of expected active information is no greater than 2^{-b}.

Thus, not only are favorable problems rare, so are favorable strategies. Whether you hold fixed the algorithm and try to match a problem to it, or hold fixed the problem and try to match a strategy to it, both are provably difficult. Because matching problems to algorithms is hard, seemingly serendipitous agreement between the two calls for further explanation, serving as evidence against blind matching by independent mechanisms (especially for very sparse targets embedded in very large spaces). More importantly, this result places hard quantitative constraints (similar to minimax bounds in statistical learning theory) on information costs for automated machine learning, which attempts to match learning algorithms to learning problems [20, 12].

5.4 Success Under Dependence

Theorem 3. *(Success Under Dependence) Define* $\tau_k = \{T \mid T \subset \Omega, |T| = k \in \mathbb{N}\}$ *and let* $\mathcal{B}_m$ *denote any set of binary strings, such that the strings are of length* m *or less. Define* q *as the expected per-query probability of success under the joint distribution on* $T \in \tau_k$ *and* $F \in \mathcal{B}_m$ *for any fixed algorithm* $\mathcal{A}$*, so that* $q := \mathbb{E}_{T,F}\left[q(T,F)\right]$*, namely,*

$$q = \mathbb{E}_{T,F}\left[\,\overline{P}(\omega \in T|F)\right] = \Pr(\omega \in T; \mathcal{A}).$$

Then,

$$q \leq \frac{I(T;F) + D(P_T \| \mathcal{U}_T) + 1}{I_\Omega}$$

where $I_\Omega = -\log k/|\Omega|$*,* $D(P_T\|\mathcal{U}_T)$ *is the Kullback-Liebler divergence between the marginal distribution on* T *and the uniform distribution on* T*, and* $I(T;F)$ *is the mutual information. Alternatively, we can write*

$$q \leq \frac{H(\mathcal{U}_T) - H(T \mid F) + 1}{I_\Omega}$$

where $H(\mathcal{U}_T) = \log \binom{|\Omega|}{k}$.

Thus the bound on expected probability of success improves monotonically with the amount of reliable information given concerning targets. We quantify the degree of reliability as the mutual information between target set T and external information resource F, under any fixed joint distribution on τ_k and $\mathcal{B}_m$. The larger the mutual information, the more information F can give us regarding the location of T, and thus, the greater the opportunity for success. We see that I_Ω measures the relative target sparseness of the search problem, and can be interpreted as the information cost of locating a target element in the absence of side-information. $D(P_T\|\mathcal{U}_T)$ is naturally interpreted as the predictability of the target sets, since large values imply the probable occurrence of only a small number of possible target sets. The mutual information $I(T;F)$ is the amount of exploitable information the external resource contains regarding T; lowering the mutual information lowers the maximum expected probability of success for the algorithm. Lastly, the 1 in the numerator upper bounds the contribution of pure randomness, as explained in the Appendix. Thus, this expression constrains the relative contributions of predictability, problem difficulty, side-information, and randomness for a successful search.

6 Examples

6.1 Landmark Security

Returning to our running example, we can now fill in a few details. Our external information resource is the pertinent intelligence data, mined through surveillance. We begin with background knowledge (represented in $F(\emptyset)$), used to make the primary security force placements. Team members on the ground communicate updates back to central headquarters, such as suspicious activity, which are the $F(\omega)$ evaluations used to update the internal information state. Each resource allocated is a query, and manpower constraints limit the number of queries available. Doing more with fewer officers is better, so the hope is to maximize the per-officer probability of stopping the attack.

Our results tell us a few things. First, a fixed strategy can only work well in a limited number of situations. There is little or no hope of a problem being a good match for your strategy if the problem arises independently of it (Theorem 1). So reliable intelligence becomes key. The better correlated the intelligence reports are with the actual plot, the better a strategy can perform (Theorem 3). However, even for a fixed search problem with reliable external information resource there is no guarantee of success, if the strategy is chosen poorly; the proportion of good strategies for a fixed problem is no better than the proportion of good problems for a fixed algorithm (Theorem 2). Thus, domain knowledge is crucial in choosing either. Without side-information to guide the match between search strategy and search problem, the expected probability of success is dismal in target-sparse situations.

6.2 One-Size-Fits-All Fitness Functions

Theorem 1 gives us an upper bound on the proportion of favorable search problems, but what happens when we have a single, fixed information resource, such as a single fitness function? A natural question to ask is for how many target locations can such a fitness function be useful. More precisely, for a given search algorithm, for what proportion of search space locations can the

fitness function raise the expected probability of success, assuming a target element happens to be located at one of those spots?

Applying Theorem 1 with $|T| = 1$, we find that a fitness function can significantly raise the probability of locating target elements placed on, at most, $1/q_{\min}$ search space elements. We see this as follows. Since $|T| = 1$ and the fitness function is fixed, each search problem maps to exactly one element of Ω, giving $|\Omega|$ possible search problems. The number of $q_{\min}$-favorable search problems is upper-bounded by

$$|\Omega|\frac{p}{q_{\min}} = |\Omega|\frac{|T|/|\Omega|}{q_{\min}} = |\Omega|\frac{1/|\Omega|}{q_{\min}} = \frac{1}{q_{\min}}.$$

Because this expression is independent of the size of the search space, the number of elements for which a fitness function can strongly raise the probability of success remains fixed even as the size of the search space increases. Thus, for very large search spaces the proportion of favored locations effectively vanishes. There can exist no single fitness function that is strongly favorable for many elements simultaneously, and thus no "one-size-fits-all" fitness function.

6.3 Binary Classification

It has been suggested that machine learning represents a type of search through parameter or hypothesis space [14]. Supporting this view, we can represent binary classification problems within our framework. The space of possible binary concepts is Ω, with the true concept being an element in that space. In our example, let $|\Omega| = 2^{100}$. The target set consists of the set of all concepts in that space that (1) are consistent with the training data (which we will assume all are), and (2) differ from the truth in at most 10% of positions on the generalization held-out dataset. Thus, $|T| = \sum_{i=0}^{10}\binom{100}{i}$. Let us assume the marginal distribution on T is uniform, which isn't necessary but simplifies the calculation. The external information resource F is the set of training examples. The algorithm uses the training examples (given by $F(\emptyset)$) to produce a distribution over the space of concepts; for deterministic algorithms, this is a degenerate distribution on exactly one element. A single query is then taken (i.e., a concept is output), and we assess the probability of success for the single query. The expected chances of outputting a concept with at least 90% generalization accuracy is thus no greater than $\frac{I(T;F)+1}{I_\Omega} \approx \frac{I(T;F)}{I_\Omega} \leq \frac{I(T;F)}{59}$. The denominator is the information cost of specifying at least one element of the target set and the numerator represents the information resources available for doing so. When the mutual information meets (or exceeds) that cost, success can be ensured for any algorithm perfectly mining the available mutual information. When noise reduces the mutual information below the information cost, the expected probability of success becomes strictly bounded in proportion to that ratio.

6.4 Hyperparameter Optimization

Given that sequential hyperparameter optimization is a literal search through a space of hyperparameter configurations, our results are directly applicable. The search space Ω consists of all the possible hyperparameter configurations (appropriately discretized in the case of numerical hyperparameters). The target set T is determined by the particular learning algorithm the configurations are applied to, the performance metric used, and the level of performance desired. Let X denote a set of points sampled from the space, and let the information gained from the sample become the external information resource f. Given that resource, we have the following theorem:

Theorem 4. *Given a search algorithm $\mathcal{A}$, a finite discrete hyperparameter configuration space Ω, a set X of points sampled from that search space, and information resource f that is a function of X, let $\Omega' := \Omega \setminus X$, $\tau_k = \{T \mid T \subset \Omega', |T| = k \in \mathbb{N}\}$, and $\tau_{k,q_{min}} = \{T \mid T \in \tau_k, q(T,f) \geq q_{min}\}$, where $q(T,f)$ is the expected per-query probability of success for algorithm $\mathcal{A}$ under T and f. Then,*

$$\frac{|\tau_{k,q_{min}}|}{|\tau_k|} \leq \frac{p'}{q_{min}}$$

where $p' = k/|\Omega'|$

The proof follows directly from Theorem 1.

The proportion of possible hyperparameter target sets giving an expected probability of success $q_{\min}$ or more is minuscule when $k \ll |\Omega'|$. If we have no additional information beyond that gained from the points X, we have no justifiable basis for expecting a successful search. Thus, we must make some assumptions concerning the relationship of the points sampled to the remaining points in Ω'. We can do so by either assuming structure on the search space, such that spatial information becomes informative, or by making an assumption on the process by which X was sampled, so that the sample is representative of the space in quantifiable ways. These assumptions allow f to become informative of the target set T, leading to exploitable dependence. Thus we see the need for inductive bias in hyperparameter optimization (to expand a term used by Mitchell [13] for classification), which hints at a strategy for creating more effective hyperparameter optimization algorithms (i.e., through exploitation of spatial structure).

6.5 Proliferation of Learning Algorithms

The Famine of Forte result also helps make sense of recent trends in machine learning. It explains why new algorithms continue to be developed, year after year, conference after conference, despite decades of research in machine learning and optimization. Given that algorithms can only perform well on a narrow subset of problems, we must either continue to devise novel algorithms for new problem domains or else move towards flexible algorithms that can modify their behavior solely though parameterization. The latter effectively behave like new strategies for new hyperparameter configurations. The explosive rise of flexible, hyperparameter sensitive algorithms like deep learning methods and vision architectures shows a definite trend towards the latter, with their hyperparameter sensitivity being well known [19, 2]. Furthermore, because flexible algorithms are highly sensitive to their hyperparameters (by definition), this explains the concurrent rise in automated hyperparameter optimization methods [18, 1, 20, 21].

7 Conclusion

A friend once remarked that the Famine of Forte results bring to mind Lake Wobegone, "where all the children are above average." In a world where every published algorithm is "above average," conservation of information reminds us that this cannot be. Improved performance over any subset of problems necessarily implies degraded performance over the remaining problems [16], and all methods have performance equivalent to random sampling when uniformly averaged over any closed-under-permutation set of problems [23, 3, 17].

But there are many ways to achieve an average. An algorithm might perform well over a large number of problems, only to perform poorly on a small set of remaining problems. A different algorithm might perform close to average on all problems, with slight variations around the mean. How large can the subset of problems with improved performance be? We show that the maximum *proportion* of search problems for which an algorithm can attain $q_{\min}$-favorable performance is bound from above by $p/q_{\min}$. Thus, an algorithm can only perform well over a narrow subset of possible problems. Not only is there no free lunch, but there is a famine of favorable search problems.

If finding a good search problem for a fixed algorithm is hard, then so is finding a good search algorithm for a fixed problem (Theorem 2). Thus, the matching of problems to algorithms is provably difficult, regardless of which is fixed and which varies.

Pessimism aside, our results become more optimistic once we restrict ourselves to those search problems for which the external information resource is strongly informative of the target set, as is often assumed to be the case. For those problems, the expected per-query probability of success is upper bounded by a function involving the mutual information between target sets and external information resources (like training datasets and objective functions). The lower the mutual information, the lower chance of success, but the bound improves as the dependence between the two is strengthened.

The search framework we propose is general enough to be applied to many problem areas, such as machine learning, evolutionary search, security and hyperparameter optimization. The results are not just of theoretical importance, but help explain real-world phenomena, such as the need for exploitable dependence in machine learning and the empirical difficulty of automated learning [12]. Our results help us understand the growing popularity of deep learning methods and unavoidable interest in automated hyperparameter tuning methods. Extending the framework to continuous settings and other problem areas (such as active learning and regression) is the focus of ongoing research.

References

[1] J. Bergstra and Y. Bengio. Random search for hyper-parameter optimization. *The Journal of Machine Learning Research*, 13(1):281–305, 2012.

[2] J. Bergstra, D. Yamins, and D. Cox. Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures. In *Proc. 30th International Conference on Machine Learning (ICML-13)*, 2013.

[3] J. Culberson. On the futility of blind search: An algorithmic view of 'no free lunch'. *Evolutionary Computation*, 6(2):109–127, 1998.

[4] W. Dembski, W. Ewert, and R. Marks II. A general theory of information cost incurred by successful search. In *Biological Information*, chapter 3, pages 26–63. World Scientific, 2013.

[5] W. Dembski and R. Marks II. Conservation of information in search: Measuring the cost of success. *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*, 39(5):1051 –1061, sept. 2009.

[6] W. Dembski and R. Marks II. The search for a search: Measuring the information cost of

higher level search. *Journal of Advanced Computational Intelligence and Intelligent Informatics*, 14(5):475–486, 2010.

[7] S. Droste, T. Jansen, and I. Wegener. Optimization with randomized search heuristics–the (a)nfl theorem, realistic scenarios, and difficult functions. *Theoretical Computer Science*, 287(1):131–144, 2002.

[8] T. English. Evaluation of evolutionary and genetic optimizers: No free lunch. In *Evolutionary Programming V: Proceedings of the Fifth Annual Conference on Evolutionary Programming*, pages 163–169, 1996.

[9] T. English. Optimization is easy and learning is hard in the typical function. In *Evolutionary Computation, 2000. Proceedings of the 2000 Congress on*, volume 2, pages 924–931. IEEE, 2000.

[10] T. English. No more lunch: Analysis of sequential search. In *Evolutionary Computation, 2004. CEC2004. Congress on*, volume 1, pages 227–234. IEEE, 2004.

[11] R. Fano and D. Hawkins. Transmission of information: A statistical theory of communications. *American Journal of Physics*, 29(11):793–794, 1961.

[12] I. Guyon, K. Bennett, G. Cawley, H. J. Escalante, S. Escalera, T. K. Ho, N. Maci, B. Ray, M. Saeed, A. Statnikov, and E. Viegas. Design of the 2015 chalearn automl challenge. In *2015 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8, July 2015.

[13] T. M. Mitchell. The need for biases in learning generalizations. Technical report, Rutgers University, 1980.

[14] T. M. Mitchell. Generalization as search. *Artificial intelligence*, 18(2):203–226, 1982.

[15] G. D. Montañez. Bounding the number of favorable functions in stochastic search. In *Evolutionary Computation (CEC), 2013 IEEE Congress on*, pages 3019–3026, June 2013.

[16] C. Schaffer. A conservation law for generalization performance. In W. W. Cohen and H. Hirsch, editors, *Proceedings of the Eleventh International Machine Learning Conference*, pages 259–265. Rutgers University, New Brunswick, NJ, 1994.

[17] C. Schumacher, M. Vose, and L. Whitley. The no free lunch and problem description length. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2001)*, pages 565–570, 2001.

[18] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. de Freitas. Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2016.

[19] J. Snoek, H. Larochelle, and R. P. Adams. Practical bayesian optimization of machine learning algorithms. In *Advances in neural information processing systems*, pages 2951–2959, 2012.

[20] C. Thornton, F. Hutter, H. H. Hoos, and K. Leyton-Brown. Auto-weka: Combined selection and hyperparameter optimization of classification algorithms. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 847–855. ACM, 2013.

[21] Z. Wang, F. Hutter, M. Zoghi, D. Matheson, and N. de Feitas. Bayesian optimization in a billion dimensions via random embeddings. *Journal of Artificial Intelligence Research (JAIR)*, 55:361–387, February 2016.

[22] D. Whitley. Functions as permutations: regarding no free lunch, walsh analysis and summary statistics. In *Parallel Problem Solving from Nature PPSN VI*, pages 169–178. Springer, 2000.

[23] D. Wolpert and W. Macready. No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation*, 1(1):67–82, April 1997.

8 Appendix: Proofs

Lemma 1. *(Expected Per Query Performance From Expected Distribution) Let t be a target set, $q(t, f)$ the expected per-query probability of success for an algorithm and ν be the conditional joint measure induced by that algorithm over finite sequences of probability distributions and search histories, conditioned on external information resource f. Denote a probability distribution sequence by $\tilde{P}$ and a search history by h. Let $\mathcal{U}(\tilde{P})$ denote a uniform distribution on elements of $\tilde{P}$ and define $\overline{P}(x \mid f) = \int \mathbb{E}_{P\sim\mathcal{U}(\tilde{P})}[P(x)]d\nu(\tilde{P}, h \mid f)$. Then,*

$$q(t, f) = \overline{P}(X \in t|f)$$

where $\overline{P}(X|f)$ is a probability distribution on the search space.

Proof. Begin by expanding the definition of $\mathbb{E}_{P\sim\mathcal{U}(\tilde{P})}[P(x)]$, being the average probability mass on element x under sequence $\tilde{P}$:

$$\mathbb{E}_{P\sim\mathcal{U}(\tilde{P})}[P(x)] = \frac{1}{|\tilde{P}|}\sum_{i=1}^{|\tilde{P}|} P_i(x).$$

Next, we confirm that $\overline{P}(x|f)$ is a proper probability distribution:

1. $\overline{P}(x|f) \geq 0$, being the integral of a nonnegative function;

2. $\overline{P}(x|f) \leq 1$, since

$$\overline{P}(x|f) \leq \int \left[\frac{1}{|\tilde{P}|}\sum_{i=1}^{|\tilde{P}|} 1\right] d\nu(\tilde{P}, h|f) = 1;$$

3. $\overline{P}(x|f)$ sums to one, since

$$\begin{aligned}
\sum_x \overline{P}(x \mid f) &= \sum_x \left[\int \frac{1}{|\tilde{P}|}\sum_{i=1}^{|\tilde{P}|} P_i(x) d\nu(\tilde{P}, h|f)\right] \\
&= \int \frac{1}{|\tilde{P}|}\sum_{i=1}^{|\tilde{P}|}\left[\sum_x P_i(x)\right] d\nu(\tilde{P}, h|f) \\
&= \int \frac{1}{|\tilde{P}|}\sum_{i=1}^{|\tilde{P}|} 1 d\nu(\tilde{P}, h|f) \\
&= \int \frac{|\tilde{P}|}{|\tilde{P}|} d\nu(\tilde{P}, h|f) \\
&= 1.
\end{aligned}$$

Finally,

$$\begin{aligned}
\overline{P}(X \in t|f) &= \sum_x \mathbb{1}_{x\in t}\overline{P}(x|f) \\
&= \sum_x \left[\mathbb{1}_{x\in t}\int \mathbb{E}_{P\sim\mathcal{U}(\tilde{P})}[P(x)] d\nu(\tilde{P}, h|f)\right] \\
&= \sum_x \left[\mathbb{1}_{x\in t}\int \left[\frac{1}{|\tilde{P}|}\sum_{i=1}^{|\tilde{P}|} P_i(x)\right] d\nu(\tilde{P}, h|f)\right] \\
&= \sum_x \int \frac{1}{|\tilde{P}|}\left[\sum_{i=1}^{|\tilde{P}|} \mathbb{1}_{x\in t} P_i(x)\right] d\nu(\tilde{P}, h|f) \\
&= \int \frac{1}{|\tilde{P}|}\sum_{i=1}^{|\tilde{P}|}\left[\sum_x \mathbb{1}_{x\in t} P_i(x)\right] d\nu(\tilde{P}, h|f) \\
&= \mathbb{E}_{\tilde{P},H}\left[\frac{1}{|\tilde{P}|}\sum_{i=1}^{|\tilde{P}|}\mathbb{E}_{X\sim P_i}[\mathbb{1}_{X\in t}]\middle| f\right] \\
&= \mathbb{E}_{\tilde{P},H}\left[\frac{1}{|\tilde{P}|}\sum_{i=1}^{|\tilde{P}|} P_i(X \in t)\middle| f\right] \\
&= q(t, f).
\end{aligned}$$

□

Lemma 2. *(Maximum Number of Satisfying Vectors) Given an integer $1 \leq k \leq n$, a set $\mathcal{S} = \{\mathbf{s} : \mathbf{s} \in \{0,1\}^n, \|\mathbf{s}\| = \sqrt{k}\}$ of all n-length k-hot binary vectors, a set $\mathcal{P} = \{P : P \in \mathbb{R}^n, \sum_j P_j = 1\}$ of discrete n-dimensional simplex vectors, and a fixed scalar threshold $\epsilon \in [0, 1]$, then for any fixed $P \in \mathcal{P}$,*

$$\sum_{\mathbf{s}\in\mathcal{S}} \mathbb{1}_{\mathbf{s}^\top P \geq \epsilon} \leq \frac{1}{\epsilon}\binom{n-1}{k-1}$$

where $\mathbf{s}^\top P$ denotes the vector dot product between $\mathbf{s}$ and P.

Proof. For $\epsilon = 0$, the bound holds trivially. For $\epsilon > 0$, let S be a random quantity that takes values $\mathbf{s}$ uniformly in the set $\mathcal{S}$. Then, for any fixed $P \in \mathcal{P}$,

$$\begin{aligned}
\sum_{\mathbf{s}\in\mathcal{S}} \mathbb{1}_{\mathbf{s}^\top P \geq \epsilon} &= \binom{n}{k}\mathbb{E}\left[\mathbb{1}_{S^\top P \geq \epsilon}\right] \\
&= \binom{n}{k}\Pr\left(S^\top P \geq \epsilon\right).
\end{aligned}$$

Let $\mathbf{1}$ denotes the all ones vector. Under a uniform distribution on random quantity S and because P does

not change with respect to $\mathbf{s}$, we have

$$\begin{aligned}
\mathbb{E}\left[S^\top P\right] &= \binom{n}{k}^{-1} \sum_{\mathbf{s}\in\mathcal{S}} \mathbf{s}^\top P \\
&= P^\top \binom{n}{k}^{-1} \sum_{\mathbf{s}\in\mathcal{S}} \mathbf{s} \\
&= P^\top \frac{\mathbf{1}\binom{n-1}{k-1}}{\binom{n}{k}} \\
&= P^\top \frac{\mathbf{1}\binom{n-1}{k-1}}{\frac{n}{k}\binom{n-1}{k-1}} \\
&= \frac{k}{n} P^\top \mathbf{1} \\
&= \frac{k}{n}
\end{aligned}$$

since P must sum to 1.

Noting that $S^\top P \geq 0$, we use Markov's inequality to get

$$\begin{aligned}
\sum_{\mathbf{s}\in\mathcal{S}} \mathbb{1}_{\mathbf{s}^\top P \geq \epsilon} &= \binom{n}{k} \Pr\left(S^\top P \geq \epsilon\right) \\
&\leq \binom{n}{k} \frac{1}{\epsilon} \mathbb{E}\left[S^\top P\right] \\
&= \binom{n}{k} \frac{1}{\epsilon} \frac{k}{n} \\
&= \frac{1}{\epsilon} \frac{k}{n} \frac{n}{k} \binom{n-1}{k-1} \\
&= \frac{1}{\epsilon} \binom{n-1}{k-1}.
\end{aligned}$$

□

Theorem 1. *(Famine of Forte) Define*

$$\tau_k = \{T \mid T \subset \Omega, |T| = k \in \mathbb{N}\}$$

and let $\mathcal{B}_m$ denote any set of binary strings, such that the strings are of length m or less. Let

$$\begin{aligned}
R &= \{(T,F) \mid T \in \tau_k, F \in \mathcal{B}_m\}, \text{ and} \\
R_{q_{min}} &= \{(T,F) \mid T \in \tau_k, F \in \mathcal{B}_m, q(T,F) \geq q_{min}\},
\end{aligned}$$

where $q(T,F)$ is the expected per-query probability of success for algorithm $\mathcal{A}$ on problem (Ω, T, F). Then for any $m \in \mathbb{N}$,

$$\frac{|R_{q_{min}}|}{|R|} \leq \frac{p}{q_{min}}$$

and

$$\lim_{m\to\infty} \frac{|R_{q_{min}}|}{|R|} \leq \frac{p}{q_{min}}$$

where $p = k/|\Omega|$.

Proof. We begin by defining a set $\mathcal{S}$ of all $|\Omega|$-length target functions with exactly k ones, namely, $\mathcal{S} = \{\mathbf{s} : \mathbf{s} \in \{0,1\}^{|\Omega|}, \|\mathbf{s}\| = \sqrt{k}\}$. For each of these, we have $|\mathcal{B}_m|$ external information resources. The total number of search problems is therefore

$$\binom{|\Omega|}{k} |\mathcal{B}_m|. \tag{3}$$

We seek to bound the proportion of possible search problems for which $q(\mathbf{s}, f) \geq q_{\min}$ for any threshold $q_{\min} \in (0,1]$. Thus,

$$\begin{aligned}
\frac{|R_{q_{\min}}|}{|R|} &\leq \frac{|\mathcal{B}_m| \sup_f \left[\sum_{\mathbf{s}\in\mathcal{S}} \mathbb{1}_{q(\mathbf{s},f)\geq q_{\min}}\right]}{|\mathcal{B}_m|\binom{|\Omega|}{k}} &(4)\\
&= \binom{|\Omega|}{k}^{-1} \sum_{\mathbf{s}\in\mathcal{S}} \mathbb{1}_{q(\mathbf{s},f^*)\geq q_{\min}}, &(5)
\end{aligned}$$

where $f^* \in \mathcal{B}_m$ denotes the arg sup of the expression. Therefore,

$$\begin{aligned}
\frac{|R_{q_{\min}}|}{|R|} &\leq \binom{|\Omega|}{k}^{-1} \sum_{\mathbf{s}\in\mathcal{S}} \mathbb{1}_{q(\mathbf{s},f^*)\geq q_{\min}} \\
&= \binom{|\Omega|}{k}^{-1} \sum_{\mathbf{s}\in\mathcal{S}} \mathbb{1}_{\overline{P}(\omega\in\mathbf{s}|f^*)\geq q_{\min}} \\
&= \binom{|\Omega|}{k}^{-1} \sum_{\mathbf{s}\in\mathcal{S}} \mathbb{1}_{\mathbf{s}^\top \overline{P}_{f^*}\geq q_{\min}}
\end{aligned}$$

where the first equality follows from Lemma 1, $\omega \in \mathbf{s}$ means the target function $\mathbf{s}$ evaluated at ω is one, and $\overline{P}_{f^*}$ represents the $|\Omega|$-length probability vector defined by $\overline{P}(\cdot|f^*)$. By Lemma 2, we have

$$\begin{aligned}
\binom{|\Omega|}{k}^{-1} \sum_{\mathbf{s}\in\mathcal{S}} \mathbb{1}_{\mathbf{s}^\top \overline{P}_{f^*}\geq q_{\min}} &\leq \binom{|\Omega|}{k}^{-1} \left[\frac{1}{q_{\min}} \binom{|\Omega|-1}{k-1}\right] \\
&= \frac{k}{|\Omega|} \frac{1}{q_{\min}} \\
&= p/q_{\min} &(6)
\end{aligned}$$

proving the result for finite external information resources.

To extend to infinite external information resources, let $A_m = \{f : f \in \{0,1\}^\ell, \ell \in \mathbb{N}, \ell \leq m\}$ and define

$$a_m := \frac{|A_m| \sup_{f\in A_m} \left[\sum_{\mathbf{s}\in\mathcal{S}} \mathbb{1}_{q(\mathbf{s},f)\geq q_{\min}}\right]}{|A_m|\binom{|\Omega|}{k}}, \tag{7}$$

$$b_m := \frac{|\mathcal{B}_m| \sup_{f\in \mathcal{B}_m} \left[\sum_{\mathbf{s}\in\mathcal{S}} \mathbb{1}_{q(\mathbf{s},f)\geq q_{\min}}\right]}{|\mathcal{B}_m|\binom{|\Omega|}{k}}. \tag{8}$$

We have shown that $A_m \leq p/q_{\min}$ for each $m \in \mathbb{N}$.

Thus,

$$\begin{aligned}\limsup_{m\to\infty}\frac{|A_m|\sup_{f\in A_m}\left[\sum_{\mathbf{s}\in\mathcal{S}}\mathbb{1}_{q(\mathbf{s},f)\geq q_{\min}}\right]}{|A_m|\binom{|\Omega|}{k}}&=\limsup_{m\to\infty}a_m\\&\leq\sup_m a_m\\&\leq p/q_{\min}.\end{aligned}$$

Next, we use the monotone convergence theorem to show the limit exists. First,

$$\lim_{m\to\infty}a_m=\lim_{m\to\infty}\frac{\sup_{f\in A_m}\left[\sum_{\mathbf{s}}\mathbb{1}_{q(\mathbf{s},f)\geq q_{\min}}\right]}{\binom{|\Omega|}{k}}\tag{9}$$

By construction, the successive A_m are nested with increasing m, so the sequence of suprema (and numerator) are increasing, though not necessarily strictly increasing. The denominator is not dependent on m, so $\{a_m\}$ is an increasing sequence. Because it is also bounded above by $p/q_{\min}$, the limit exists by monotone convergence. Thus,

$$\lim_{m\to\infty}a_m=\limsup_{m\to\infty}a_m\leq p/q_{\min}.$$

Lastly,

$$\begin{aligned}\lim_{m\to\infty}b_m&=\lim_{m\to\infty}\frac{|\mathcal{B}_m|\sup_{f\in\mathcal{B}_m}\left[\sum_{\mathbf{s}\in\mathcal{S}}\mathbb{1}_{q(\mathbf{s},f)\geq q_{\min}}\right]}{|\mathcal{B}_m|\binom{|\Omega|}{k}}\\&=\lim_{m\to\infty}\frac{\sup_{f\in\mathcal{B}_m}\left[\sum_{\mathbf{s}\in\mathcal{S}}\mathbb{1}_{q(\mathbf{s},f)\geq q_{\min}}\right]}{\binom{|\Omega|}{k}}\\&\leq\lim_{m\to\infty}\frac{\sup_{f\in A_m}\left[\sum_{\mathbf{s}\in\mathcal{S}}\mathbb{1}_{q(\mathbf{s},f)\geq q_{\min}}\right]}{\binom{|\Omega|}{k}}\\&=\lim_{m\to\infty}a_m\\&\leq p/q_{\min}.\end{aligned}$$

□

Corollary 1. *(Conservation of Expected Active Information) Define* active information *as $I_+=-\log p/q$, where p is the per-query probability of success for uniform random sampling and q is the probability of success for an alternative search algorithm. Define*

$$\tau_k=\{T\mid T\subset\Omega,|T|=k\in\mathbb{N}\}$$

and let $\mathcal{B}_m$ denote any set of binary strings, such that the strings are of length m or less. Let

$$\begin{aligned}R&=\{(T,F)\mid T\in\tau_k,F\in\mathcal{B}_m\},\ \textit{and}\\R_b&=\{(T,F)\mid T\in\tau_k,F\in\mathcal{B}_m,\mathbb{E}[I_+]\geq b\}.\end{aligned}$$

Then for any $m\in\mathbb{N}$

$$\frac{|R_b|}{|R|}\leq 2^{-b}.$$

Proof. The proof follows from the definition of active information, Jensen's inequality and Theorem 1. Let

$$I_+=-\log\left(\frac{p}{q}\right),\tag{10}$$

$$b\leq\mathbb{E}\left[-\log\left(\frac{p}{q}\right)\right].\tag{11}$$

By Jensen's inequality and the concavity of $\log(q/p)$ in q, we have

$$\begin{aligned}b&\leq\mathbb{E}\left[-\log\left(\frac{p}{q}\right)\right]\\&=\mathbb{E}\left[\log\left(\frac{q}{p}\right)\right]\\&\leq\log\left(\frac{\mathbb{E}\left[q\right]}{p}\right)\\&=-\log\left(\frac{p}{q_{\min}}\right),\end{aligned}$$

where we have defined $q_{\min}:=\mathbb{E}\left[q\right]$. Rearranging the first and last expressions in the series of inequalities yields

$$q_{\min}\geq p2^b.$$

Since $\mathbb{E}[I_+]\geq b$ implies $q_{\min}\geq p2^b$, the set of problems for which $\mathbb{E}[I_+]\geq b$ can be no bigger than the set for which $q_{\min}\geq p2^b$.

By Theorem 1, the proportion of problems for which $q_{\min}$ is at least $p2^b$ is no greater than $p/(p2^b)$. Thus,

$$\frac{|R_b|}{|R|}\leq\frac{1}{2^b}.\tag{12}$$

□

Lemma 3. *(Maximum Proportion of Satisfying Strategies) Given an integer $1\leq k\leq n$, a set $\mathcal{S}=\{\mathbf{s}:\mathbf{s}\in\{0,1\}^n,\|\mathbf{s}\|=\sqrt{k}\}$ of all n-length k-hot binary vectors, a set $\mathcal{P}=\{P:P\in\mathbb{R}^n,\sum_j P_j=1\}$ of discrete n-dimensional simplex vectors, and a fixed scalar threshold $\epsilon\in[0,1]$, then*

$$\max_{\mathbf{s}\in\mathcal{S}}\frac{\mu(\mathcal{G}_{\mathbf{s},\epsilon})}{\mu(\mathcal{P})}\leq\frac{1}{\epsilon}\frac{k}{n}$$

where $\mathcal{G}_{\mathbf{s},\epsilon}=\{P:P\in\mathcal{P},\mathbf{s}^\top P\geq\epsilon\}$ and μ is Lebesgue measure.

Proof. Similar results have been proved by others with regard to No Free Lunch theorems [23, 5, 17, 4, 8]. Our result concerns the maximum proportion of sufficiently good strategies (not the mean performance of strategies over all problems, as in the NFL case) and is a simplification over previous search-for-a-search results.

For $\epsilon = 0$, the bound holds trivially. For $\epsilon > 0$, We first notice that the $\mu(\mathcal{P})^{-1}$ term can be viewed as a uniform density over the region of the simplex $\mathcal{P}$, so that the integral becomes an expectation with respect to this distribution, where P is drawn uniformly from $\mathcal{P}$. Thus, for any $\mathbf{s} \in \mathcal{S}$,

$$\begin{aligned}\frac{\mu(\mathcal{G}_{\mathbf{s},\epsilon})}{\mu(\mathcal{P})} &= \int_{\mathcal{P}} \frac{1}{\mu(\mathcal{P})} \left[\mathbb{1}_{\mathbf{s}^\top P \geq \epsilon}\right] d\mu(P) \\ &= \mathbb{E}_{P\sim\mathcal{U}(\mathcal{P})}\left[\mathbb{1}_{\mathbf{s}^\top P \geq \epsilon}\right] \\ &= \Pr(\mathbf{s}^\top P \geq \epsilon) \\ &\leq \frac{1}{\epsilon}\mathbb{E}_{P\sim\mathcal{U}(\mathcal{P})}\left[\mathbf{s}^\top P\right],\end{aligned}$$

where the final line follows from Markov's inequality. Since the symmetric Dirichlet distribution in n dimensions with parameter $\alpha = 1$ gives the uniform distribution over the simplex, we get

$$\begin{aligned}\mathbb{E}_{P\sim\mathcal{U}(\mathcal{P})}[P] &= \mathbb{E}_{P\sim\mathrm{Dir}(\alpha=1)}[P] \\ &= \left(\frac{\alpha}{\sum_{i=1}^n \alpha}\right)\mathbf{1} \\ &= \left(\frac{1}{n}\right)\mathbf{1},\end{aligned}$$

where $\mathbf{1}$ denotes the all ones vector. We have

$$\begin{aligned}\frac{1}{\epsilon}\mathbb{E}_{P\sim\mathcal{U}(\mathcal{P})}\left[\mathbf{s}^\top P\right] &= \frac{1}{\epsilon}\mathbf{s}^\top \mathbb{E}_{P\sim\mathcal{U}(\mathcal{P})}[P] \\ &= \frac{1}{\epsilon}\mathbf{s}^\top\left(\frac{1}{n}\right)\mathbf{1} \\ &= \frac{1}{\epsilon}\frac{k}{n}.\end{aligned}$$

□

Theorem 2. *(Famine of Favorable Strategies) For any fixed search problem (Ω, t, f), set of probability mass functions $\mathcal{P} = \{P : P \in \mathbb{R}^{|\Omega|}, \sum_j P_j = 1\}$, and a fixed threshold $q_{min} \in [0,1]$,*

$$\frac{\mu(\mathcal{G}_{t,q_{min}})}{\mu(\mathcal{P})} \leq \frac{p}{q_{min}},$$

where $\mathcal{G}_{t,q_{min}} = \{P : P \in \mathcal{P}, t^\top P \geq q_{min}\}$ and μ is Lebesgue measure. Furthermore, the proportion of possible search strategies giving at least b bits of expected active information is no greater than 2^{-b}.

Proof. Applying Lemma 3, with $\mathbf{s} = t$, $\epsilon = q_{\min}$, $k = |t|$, $n = |\Omega|$, and $p = \frac{|t|}{|\Omega|}$, yields the first result, while following the same steps as Corollary 1 gives the second. □

Theorem 3. *(Probability of Success Under Dependence) Define $\tau_k = \{T \mid T \subset \Omega, |T| = k \in \mathbb{N}\}$ and let $\mathcal{B}_m$ denote any set of binary strings, such that the strings are of length m or less. Define q as the expected per-query probability of success under the joint distribution on $T \in \tau_k$ and $F \in \mathcal{B}_m$ for any fixed algorithm $\mathcal{A}$, so that $q := \mathbb{E}_{T,F}[q(T,F)]$, namely,*

$$q = \mathbb{E}_{T,F}\left[\,\overline{P}(\omega \in T|F)\right] = \Pr(\omega \in T; \mathcal{A}).$$

Then,

$$q \leq \frac{I(T;F) + D(P_T\|\mathcal{U}_T) + 1}{I_\Omega}$$

where $I_\Omega = -\log k/|\Omega|$, $D(P_T\|\mathcal{U}_T)$ is the Kullback-Liebler divergence between the marginal distribution on T and the uniform distribution on T, and $I(T;F)$ is the mutual information. Alternatively, we can write

$$q \leq \frac{H(\mathcal{U}_T) - H(T \mid F) + 1}{I_\Omega}$$

where $H(\mathcal{U}_T) = \log\binom{|\Omega|}{k}$.

Proof. This proof loosely follows that of Fano's Inequality [11], being a reversed generalization of it, so in keeping with the traditional notation we let $\widehat{X} := \omega$ for the remainder of this proof. Let $Z = \mathbb{1}(\widehat{X} \in T)$. Using the chain rule for entropy to expand $H(Z,T|\widehat{X})$ in two different ways, we get

$$\begin{aligned}H(Z,T|\widehat{X}) &= H(Z|T,\widehat{X}) + H(T|\widehat{X}) &(13)\\ &= H(T|Z,\widehat{X}) + H(Z|\widehat{X}). &(14)\end{aligned}$$

By definition, $H(Z|T,\widehat{X}) = 0$, and by the data processing inequality $H(T|F) \leq H(T|\widehat{X})$. Thus,

$$H(T|F) \leq H(T|Z,\widehat{X}) + H(Z|\widehat{X}). \quad (15)$$

Define $P_g = \Pr(\widehat{X} \in T; \mathcal{A}) = \Pr(Z = 1)$. Then,

$$\begin{aligned}H(T|Z,\widehat{X}) &= (1-P_g)H(T|Z=0,\widehat{X}) + P_g H(T|Z=1,\widehat{X}) &(16)\\ &\leq (1-P_g)\log\binom{|\Omega|}{k} + P_g\log\binom{|\Omega|-1}{k-1} &(17)\\ &= \log\binom{|\Omega|}{k} - P_g\log\frac{|\Omega|}{k}. &(18)\end{aligned}$$

We let $H(\mathcal{U}_T) = \log\binom{|\Omega|}{k}$, being the entropy of the uniform distribution over k-sparse target sets in Ω. Therefore,

$$H(T|F) \leq H(\mathcal{U}_T) - P_g\log\frac{|\Omega|}{k} + H(Z|\widehat{X}). \quad (19)$$

Using the definitions of conditional entropy and I_Ω, we get

$$H(T) - I(T;F) \leq H(\mathcal{U}_T) - P_g I_\Omega + H(Z|\widehat{X}), \quad (20)$$

which implies

$$P_g I_\Omega \leq I(T;F) + H(\mathcal{U}_T) - H(T) + H(Z|\widehat{X}) \quad (21)$$

$$= I(T;F) + D(P_T \| \mathcal{U}_T) + H(Z|\widehat{X}). \quad (22)$$

Examining $H(Z|\widehat{X})$, we see it captures how much entropy of Z is due to the randomness of T. To see this, imagine Ω is a roulette wheel and we place our bet on $\widehat{X}$. Target elements are "chosen" as balls land on random slots, according to the distribution on T. When a ball lands on $\widehat{X}$ as often as not (roughly half the time), this quantity is maximized. Thus, this entropy captures the contribution of dumb luck, being averaged over all $\widehat{X}$. (When balls move towards always landing on $\widehat{X}$, something other than luck is at work.) We upperbound this by its maximum value of 1 and obtain

$$\Pr(\widehat{X} \in T; \mathcal{A}) \leq \frac{I(T;F) + D(P_T \| \mathcal{U}_T) + 1}{I_\Omega}, \quad (23)$$

and substitute q for $\Pr(\widehat{X} \in T; \mathcal{A})$ to obtain the first result, noting that $q = \mathbb{E}_{T,F}\left[\, \overline{P}(\omega \in T|F)\right]$ specifies a proper probability distribution by the linearity and boundedness of the expectation. To obtain the second form, use the definitions $I(T;F) = H(T) - H(T|F)$ and $D(P_T \| \mathcal{U}_T) = H(\mathcal{U}_T) - H(T)$. □